arXiv:1605.08328v1 [stat.ME] 24 May 2016

Discussion of Peters, Bühlmann and Meinshausen

Chris. J. Oates[1,2], Jessica Kasza[3,4] and Sach Mukherjee[5]
[1]University of Technology Sydney
[2]ARC Centre of Excellence for Mathematical and Statistical Frontiers
[3]Monash University
[4]Victorian Centre for Biostatistics
[5]German Center for Neurodegenerative Diseases, Bonn.

May 27, 2016

Abstract

Contribution to the discussion of the paper "Causal inference using invariant prediction: identification and confidence intervals" by Peters, Bühlmann and Meinshausen, to appear in the Journal of the Royal Statistical Society, Series B.

The information theorist is taught to list invariances and then derive models that exhibit those invariances. We warmly congratulate the authors on their insight and creativity in bringing ideas of invariance to bear on causal inference and, in so doing, furthering our understanding of the connection between causality and prediction, two areas that are often treated as separate (if not opposed!) but that are in fact closely related.

An attractive feature of the proposed method is that it allows integration of multiple data sources in a natural way, even when the precise nature or target of perturbations to the data are unknown. On the other hand, in some cases a softer approach that allows for variation in causal structure across data subsets might be appropriate (see e.g Oates *et al.*, 2014).

Returning to invariance, we wonder how far we can push the information theorist. For instance, consider estimation of the causal effect θ_{ij} of one variable X_i on another X_j. A correct causal graph G can be interrogated to produce a minimal sufficient set S of variables to adjust for when estimating θ_{ij} (Pearl, 2009). These variables S can then be included in a propensity score model, leading ultimately to an estimate $\hat{\theta}_{ij}(S)$ for θ_{ij}. Often such minimal sufficient adjustment sets are not unique, in which case any other minimal sufficient adjustment set S' will do (in the sense of allowing consistent estimation). Then, we would expect, for large sample sizes, $\hat{\theta}_{ij}(S) \approx \hat{\theta}_{ij}(S')$. On the other hand, there seems no particular reason to expect that these two estimates would coincide if the graph G were incorrect. This seems to suggest another invariance that could be exploited for causal discovery. Potentially,

other invariances could play a role. This may in future lead to having to ask which invariances are most useful in practice.

There has long been (in our view justifiable) empirical skepticism towards *de novo* causal discovery, a skepticism expressed by, among others, David Freedman and Paul Humphreys in a paper titled "Are there algorithms that discover causal structure?" (Freedman and Humphreys, 1999). The issue is that it is difficult to empirically validate causal discovery on a given problem using data at hand, leaving the analyst unsure as to whether or not the output of a given procedure should be trusted. This goes a bit further than familiar issues of statistical uncertainty, since the underlying concern is of a potentially profound mismatch between critical assumptions and the real data-generating system. The authors' insightful discussion of model mis-specification is therefore welcome and the conservative behaviour of their procedure very appealing. We note also that background scientific knowledge may itself be mis-specified but that in some circumstances it may be possible to effect "repairs" on the relevant causal structures (Oates *et al.*, 2016). We see it as a very positive development that empirical validation of causal discovery is becoming more common (see e.g. Hill *et al.*, 2016). In the near future, empirical work, not least in biology, ought to give us a better sense of the practical efficacy of causal discovery. A concrete answer to Freedman and Humphreys' question may then start to come into reach.

References

Freedman, D. and Humphreys, P. (1999) Are there algorithms that discover causal structure? *Synthese*, **121**(1-2), 29-54.

Hill, S.M. *et al.* (2016) Inferring causal molecular networks: empirical assessment through a community-based effort. *Nature Methods*, **13**(4), 310-318.

Oates, C.J., Korkola, J., Gray, J.W., Mukherjee, S. (2014) Joint Estimation of Multiple Related Biological Networks. *The Annals of Applied Statistics*, **8**(3):1892-1919.

Oates, C.J., Kasza, J. Simpson, J.A., Forbes, A.B. (2016) A pre-processing approach to repair of misspecified causal diagrams. In submission.

Pearl, J. (2009) *Causality*. Cambridge University Press.